Abstract.

We extend the Golub-Kahan algorithm for computing the singular value decomposition of bidiagonal matrices to triangular matrices. Instead of forming $R^T R$ or RR^T, our algorithm starts with the original matrix R, and generates new iterates by repeated QR factorisations.

The results of the convergence analysis of our algorithm also apply to the QR algorithm for computing eigenvalues of symmetric matrices. Our analysis demonstrates monotonic convergence of singular values and singular vectors, and a convergence rate for singular values that equals the square of the convergence rate for singular vectors. It is also possible to explain the occurrence of deflation in the interior of the matrix.

We describe the relationship between our algorithm, and the algorithms for rank revealing QR and URV decompositions. As a consequence we obtain new algorithms for computing URV decompositions, and a divide-and-conquer algorithm that computes singular values of dense matrices and may be beneficial on a parallel architecture.

We present a simple deflation and convergence criterion for triangular matrices that recognises convergence of the singular values earlier than the traditional perturbation bounds. In particular, it allows high relative accuracy in the smallest singular value.

Analysis of a QR Algorithm for Computing Singular Values

Shivkumar Chandrasekaran, Ilse Ipsen

Research Report YALEU/DCS/RR-917
August 1992

The work presented in this paper was supported by NSF grant CCR-9102853.

1 Introduction

We present a new algorithm for computing the singular value decomposition of a real upper triangular matrix R. It represents an extension to triangular matrices of the algorithm for computing singular values of a bi-diagonal matrix by Golub and Kahan [15]. Our algorithm, first proposed in [20], proceeds by determining the new iterate from a QR factorisation of the transpose of the old iterate,

$$R^{(0)} = R, \qquad {R^{(i)}}^T = Q^{(i+1)} R^{(i+1)},$$

and so avoids the explicit formation of $R^T R$ or RR^T. Fernando and Parlett [12] recently used this iteration to derive a version of Rutishauser's differential QD algorithm that computes singular values of bi-diagonal matrices to high relative accuracy. Although it would be possible to incorporate shifts, as in [12] for example, we only discuss the unshifted version in this paper.

The repeated transformation from lower triangular form ${R^{(i)}}^T$ to upper triangular form $R^{(i+1)}$ by means of orthogonal transformations $Q^{(i)}$ was motivated by an algorithm for computing partial correlation coefficients of a matrix R [9, 10] without the explicit formation of $R^T R$ or RR^T.

Although we do not, at this point, advocate the above algorithm as a practical method for computing singular values of dense matrices its analysis provides more insight into the behaviour of the Golub-Kahan algorithm for bi-diagonal matrices [15], and into the unshifted QR algorithm [18, 22, 30] for computing eigenvalues of symmetric matrices. Section 2 exposes the relation of our new algorithm to the QR and Cholesky LR algorithms for computing eigenvalues.

The convergence properties of the unshifted QR algorithm are well-known [18, 22, 23, 27, 29, 30]. They are usually derived from the fact that one iteration of the QR algorithm is mathematically equivalent to one nested subspace iteration, applied to particular starting spaces, cf. in particular [23, 27, 29]. The subspace iterates converge linearly to eigenspaces with an asymptotic convergence rate equal to a ratio of adjacent singular values (in fact, the distance between the iterates and the eigenspace decreases from the start [29]). In [28] these results are extended to the computation of the SVD of R from $R^T R$ and RR^T. The monotonic convergence of the eigenvalues during nested subspace iteration is proved through the connection to Toda flows [21].

In Section 3, we provide a simple convergence analysis of our algorithm that emphasizes the monotonic aspects of the convergence. The additional insight we get from our analysis is obtained by studying the structure of the SVD of triangular matrices. In particular, we show the monotonic decrease of the tangent of the angle between certain canonical spaces and the invariant subspaces of the iterates $R^{(i)}$, with the usual convergence rate. As for the singular values, their convergence rate is equal to the square of that of the singular vectors. These results explain the occurrence of deflation in the interior of the matrix.

Our analysis also makes it possible to better understand the relation between algorithms that produce complete SVDs and those that produce incomplete SVDs, such as algorithms for rank-revealing QR decompositions [7] and for URV decompositions [25]. In Section 4 we show that from the point of view of an individual singular value σ our algorithm proceeds in two phases: a rank-revealing phase where the singular values are separated into two groups according to whether they are larger or smaller than σ; and a monotonic phase, where the iterates converge monotonically to block-diagonal form.

Hence, preceding our algorithm with a rank-revealing algorithm accomplishes two things: it reverses the grading of inappropriately graded matrices and so enhances subsequent convergence; and it forces premature deflation of a particular off-diagonal block and thus amounts to the computation of a URV decomposition. According to this last observation we sketch a divide-and-conquer algorithm for computing singular values of dense matrices, which may be advantageous on a parallel architecture.

In Section 5 we present a very simple deflation and convergence criterion for triangular matrices that recognises convergence of the singular values earlier than the traditional perturbation bounds. In particular, it allows high relative accuracy in the smallest singular value. Our criterion represents a generalisation of the convergence criteria for bi-diagonal matrices in [11].

Section 6 concludes with a summary of the results in this paper.

Notation

We use $\|\cdot\|$ to represent the Euclidean two-norm. The identity matrix of order k is denoted by I_k and its ith column by e_i.

2 The Algorithm

In 1965 Golub and Kahan [15] introduced an algorithm for the computation of the singular values and vectors of a real upper bi-diagonal matrix B. The algorithm is based on the QR algorithm and involves only an implicit formation of the tridiagonal matrix $B^T B$. An Algol implementation of this algorithm was proposed by Golub and Reinsch in 1970 [17]. In this paper we extend the Golub-Kahan algorithm from bidiagonal matrices to triangular matrices, as was first proposed in [20].

We start with the unshifted QR algorithm for computing eigenvalues [18, 22, 30]. Given a real symmetric matrix A, the QR algorithm first determines a QR decomposition $A = QR$, where Q is orthogonal and R is upper triangular, and then it forms $\hat{A} = RQ$. The new iterate $\hat{A}$ is orthogonally similar to A since $\hat{A} = Q^T AQ$. Hence $\hat{A}$ has the same eigenvalues as A.

Our algorithm computes the singular values of a real upper triangular matrix R_0. In the first iteration, the QR decomposition of the lower triangular matrix $R_0^T = Q_1 R_1$ is determined. In the second iteration, the transpose of the resulting upper triangular matrix R_1 is in turn decomposed $R_1^T = Q_2 R_2$. Then

$$R_2 = Q_2^T R_1^T = Q_2^T R_0 Q_1,$$

and the second iterate R_2 is related to the original matrix R_0 by an orthogonal equivalence transformation. Hence R_2 has the same singular values as R_0.

2.1 Relation to QR Algorithm

The transformations

$$R_0^T = Q_1 R_1, \qquad R_1^T = Q_2 R_2$$

are mathematically equivalent to one step of the unshifted QR algorithm applied to $R_0^T R_0$. The first transformation corresponds to a QR decomposition $R_0^T R_0 = Q_1 R$, where Q_1 is orthogonal and $R = R_1 R_0$ is upper triangular. The second transformation corresponds to the completion of the similarity transformation involving Q_1,

$$R_2^T R_2 = R_1 R_1^T = Q_1^T \left(R_0^T R_0\right) Q_1.$$

If R_0 is upper bi-diagonal, so are R_1 and R_2, and the two transformations amount to applying one iteration of the Golub-Kahan algorithm [15] to R_0.

Moreover, the transformations also amount to one step of the unshifted QR algorithm [18, 22, 30] applied to $R_0 R_0^T$ because $R_0 R_0^T = Q_2(R_2 R_1)$ represents a QR decomposition of $R_0 R_0^T$, and

$$R_2 R_2^T = R_1^T R_1 = Q_2^T \left(R_0 R_0^T\right) Q_2^T$$

is the corresponding similarity transformation.

2.2 Relation to Cholesky Factorisation and LR Algorithm

If R_0 is the factor from the upper-lower Cholesky factorisation of a symmetric positive semi-definite matrix A, $A = R_0 R_0^T$, then also $A = R_1^T R_1$. So, R_1 is the factor from the lower-upper Cholesky factorisation of A, and the two factors are related through Q_1. The fact that the two Cholesky factors of a matrix are related by an orthogonal transformation is used in [9, 10] to compute partial correlation coefficients.

It is observed in [12] that the transformation $R_0^T = Q_1 R_1$ is mathematically equivalent to one iteration of the Cholesky LR algorithm [30] applied to $A_0 = R_0 R_0^T$: factor $A_0 = R_1^T R_1$, and multiply in reverse order $A_1 = R_1 R_1^T$. A second iteration of the Cholesky LR algorithm factors $A_1 = R_2^T R_2$ and multiplies $A_2 = R_2 R_2^T$. Hence the known result, Section 8.51 in [30], that one iteration of the QR algorithm for symmetric matrices is mathematically equivalent to two iterations of the LR Cholesky algorithm.

2.3 Summary

Given a real upper triangular matrix R_0, the iterations

$$R_i^T = Q_{i+1} R_{i+1}, \qquad i \geq 0$$

represent the extension of the Golub-Kahan algorithm from bi-diagonal to triangular matrices. Two iterations of this extended algorithm correspond to the implicit application of one QR iteration to $R_0 R_0^T$ and $R_0^T R_0$ because

$$R_{i+2} R_{i+2}^T = Q_{i+2}^T \left(R_i R_i^T\right) Q_{i+2}, \qquad R_{i+2}^T R_{i+2} = Q_{i+1}^T \left(R_i^T R_i\right) Q_{i+1}.$$

The above iteration represents one algorithm in a larger class of algorithms analysed in [28, 29]. There the iterates for a matrix A are obtained by applying the QR algorithm to the explicitly formed matrices $A^T A$ and $A A^T$.

3 Monotonic Convergence Results

In this section we determine some of the quantities that undergo monotonic changes during the extended Golub-Kahan algorithm.

One iteration of the Golub-Kahan algorithm applied to a real upper triangular matrix R of order n determines the decomposition $R^T = Q\hat{R}$. Partition the matrices so as to distinguish a (1,1) block of order k,

$$R = \begin{pmatrix} R_{11} & R_{12} \\ & R_{22} \end{pmatrix}, \qquad Q = \begin{pmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{pmatrix}, \qquad \hat{R} = \begin{pmatrix} \hat{R}_{11} & \hat{R}_{12} \\ & \hat{R}_{22} \end{pmatrix}.$$

We assume that the singular values of R are sorted in descending order,

$$\sigma_1 \geq \ldots \geq \sigma_k \geq \sigma_{k+1} \geq \ldots \geq \sigma_n.$$

3.1 Zero Singular Values

The following analysis requires that the matrix R be non-singular. If R were singular, one could precede our algorithm by a QR decomposition with column pivoting to enforce the disclosure of the zero singular values as follows.

First use Golub's algorithm [4, 14, 18] to perform a QR decomposition with column pivoting on the given matrix R in order to move the zeros to the bottom of the matrix. This yields

$$R = Q_P \begin{pmatrix} R_{11} & R_{12} \\ 0 & 0 \end{pmatrix} P^T,$$

where Q_P is an orthogonal matrix, P is a permutation matrix and R_{11} is a non-singular upper triangular matrix. Then perform a QR decomposition on the transpose of the resulting triangular matrix in order to eliminate the off-diagonal block,

$$\begin{pmatrix} R_{11}^T & 0 \\ R_{12}^T & 0 \end{pmatrix} = Q \begin{pmatrix} \hat{R}_{11} & 0 \\ 0 & 0 \end{pmatrix}.$$

Our algorithm can now be continued on the non-singular triangular matrix $\hat{R}_{11}$.

In exact arithmetic it therefore takes two QR decompositions to extract the zero singular values from a triangular matrix. Note that one iteration of the QR algorithm is necessary in order to expose the zero eigenvalues of an unreduced symmetric tri-diagonal matrix, Section 8.8 in [22]. Consequently, we may assume from now on that R is non-singular.

3.2 Convergence of Singular Values

From $R^T = Q\hat{R}$ and $\hat{R} = Q^T R^T$ it follows that

$$R_{11}^T = Q_{11}\hat{R}_{11}, \qquad \hat{R}_{22} = Q_{22}^T R_{22}^T,$$

so

$$\|\hat{R}_{11}^{-1}\| \leq \|R_{11}^{-1}\|, \qquad \|\hat{R}_{22}\| \leq \|R_{22}\|,$$

where $\|\cdot\|$ denotes the two-norm. These inequalities are special cases of the monotonicity properties of eigenvalues during subspace iteration [21]).

Hence

$$\|R_{11}^{(i+1)^{-1}}\| \leq \|R_{11}^{(i)^{-1}}\|, \qquad \|R_{22}^{(i+1)}\| \leq \|R_{22}^{(i)}\|,$$

and $\|R_{11}^{(i+1)^{-1}}\|$ and $\|R_{22}^{(i+1)}\|$ are monotone non-increasing sequences. Their convergence rate will be determined in Section 3.6.

3.3 SVD of Triangular Matrices

In order to understand why our algorithm makes progress in every iteration we first study the structure of the singular value decompositon of triangular matrices.

Let $R = U\Sigma V^T$ be the singular value decomposition (SVD) of R, where U and V are orthogonal matrices of order n, and Σ is a diagonal matrix, whose diagonal contains the singular values in descending order,

$$\Sigma = \begin{pmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_n \end{pmatrix}, \qquad \sigma_1 \geq \ldots \geq \sigma_k \geq \sigma_{k+1} \geq \ldots \geq \sigma_n.$$

Partition the matrices in the SVD conformally with R,

$$U = \begin{pmatrix} U_{11} & U_{12} \\ U_{21} & U_{22} \end{pmatrix}, \qquad V = \begin{pmatrix} V_{11} & V_{12} \\ V_{21} & V_{22} \end{pmatrix}, \qquad \Sigma = \begin{pmatrix} \Sigma_1 & \\ & \Sigma_2 \end{pmatrix},$$

where

$$\Sigma_1 = \begin{pmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_k \end{pmatrix}, \qquad \Sigma_2 = \begin{pmatrix} \sigma_{k+1} & & \\ & \ddots & \\ & & \sigma_n \end{pmatrix}.$$

The fact that R is triangular results in the following relations. From $RV = U\Sigma$ one gets

$$R_{22}V_{21} = U_{21}\Sigma_1, \qquad R_{22}V_{22} = U_{22}\Sigma_2,$$

and from $U^T R = \Sigma V^T$,

$$U_{11}^T R_{11} = \Sigma_1 V_{11}^T, \qquad U_{12}^T R_{11} = \Sigma_2 V_{12}^T.$$

Furthermore, $R_{21} = 0$ in $R = U\Sigma V^T$ yields

$$U_{21}\Sigma_1 V_{11}^T + U_{22}\Sigma_2 V_{12}^T = 0.$$

If V_{11} is non-singular then the last equality can be written as

$$U_{22}^{-1}U_{21} = -\Sigma_2 V_{12}^T V_{11}^{-T}\Sigma_1^{-1}.$$

In order to derive a geometric interpretation for this equality, we first digress to establish some properties of orthgonal matrices.

The CS decomposition, Theorem 2.6.1 in [18], of an orthogonal matrix

$$Z \;=\; \begin{matrix} & k & \\ k & \begin{pmatrix} Z_{11} & Z_{12} \\ Z_{21} & Z_{22} \end{pmatrix}, \end{matrix}$$

implies for non-singular Z_{11}

$$\|Z_{12}\| = \|Z_{21}\| = \sin\theta, \qquad \|Z_{11}^{-1}Z_{12}\| = \|Z_{22}^{-1}Z_{21}\| = \tan\theta,$$

where θ is some angle, and

$$\sqrt{1 - \frac{1}{\|Z_{11}^{-1}\|^2}} = \|Z_{21}\|.$$

These equalities give, together with $\sigma_{min}(Z_{11}) = 1/\|Z_{11}^{-1}\|$,

$$\sin\theta = \sqrt{1 - \sigma_{min}^2(Z_{11}^T I_k)}.$$

But the square-root term represents the distance between the column space of $\begin{pmatrix} Z_{11} \\ Z_{21} \end{pmatrix}$ and the canonical space $\begin{pmatrix} I_k \\ 0 \end{pmatrix}$, Corollary 2.6.2 in [18]. Hence, the angle θ from the CS decomposition of Z must be the largest principal angle, Section 12.4.3 in [18], between these two subspaces.

Applying this to the above inequality relating the two singular vector matrices of R yields

$$\tan\theta_{u,k} \le \frac{\sigma_{k+1}}{\sigma_k} \tan\theta_{v,k},$$

where $\theta_{u,k}$ is the largest principal angle between the leading k columns of U and the canonical subspace $\begin{pmatrix} I_k \\ 0 \end{pmatrix}$; while $\theta_{v,k}$ is the analogous angle for V. This means, the leading k columns of the left singular vector matrix U of R are closer to *canonical form* than those of the right singular vector matrix V by a factor of σ_{k+1}/σ_k. By 'canonical form' we mean an orthogonal matrix whose columns span the canonical space $\begin{pmatrix} I_k \\ 0 \end{pmatrix}$, that is, a matrix of the form $\begin{pmatrix} Z \\ 0 \end{pmatrix}$ where Z is orthogonal.

We sum up the relationship between left and right singular vector matrices of a triangular matrix in the following theorem.

Theorem 3.3.1 *Let R be a non-singular upper triangular matrix of order n with SVD $R = U\Sigma V^T$ and singular values*

$$\sigma_1 \ge \ldots \ge \sigma_k \ge \sigma_{k+1} \ge \ldots \ge \sigma_n.$$

If the leading principal submatrices of order k of U and V are non-singular then

$$\tan\theta_{u,k} \le \frac{\sigma_{k+1}}{\sigma_k} \tan\theta_{v,k},$$

where $\theta_{v,k}$ is the angle between the leading k columns of V and the canonical subspace $\begin{pmatrix} I_k \\ 0 \end{pmatrix}$; and $\theta_{u,k}$ is the corresponding angle for U.

This means, if the singular vector matrices of a non-singular upper triangular matrix R are strongly non-singular and if its singular values are well-separated then the left singular vector matrix is closer to canonical form than the right singular vector matrix (a matrix is called 'strongly non-singular' if all its leading principal submatrices are non-singular).

3.4 The Effect of One QR Decomposition on the SVD

If $R = U\Sigma V^T$ is the SVD of the original matrix then the SVD of the matrix $\hat{R}$ resulting from a QR decomposition $R^T = Q\hat{R}$ is given by

$$\hat{R} = \hat{U}\Sigma\hat{V}^T, \qquad \hat{U} = Q^T V, \quad \hat{V} = U.$$

This means a QR decomposition of R^T puts the left singular vector U matrix in place of the right singular vector matrix and creates a new left singular vector matrix $\hat{U}$. With the results from the previous section we can now see the effect of one QR decomposition on the singular vector matrices.

If the right singular vector matrix V of the original matrix R has a non-singular leading principal submatrix V_{11} of order k then

$$\tan\theta_{u,k} \leq \frac{\sigma_{k+1}}{\sigma_k} \tan\theta_{v,k},$$

where $\theta_{u,k}$ and $\theta_{v,k}$ are the respective angles between the leading k columns of U and V with the canonical subspace $\begin{pmatrix} I_k \\ 0 \end{pmatrix}$. But since U is the new right singular vector matrix $\hat{V}$, this means

$$\tan\theta_{\hat{v},k} \leq \frac{\sigma_{k+1}}{\sigma_k} \tan\theta_{v,k}.$$

Moreover, the matrix $\hat{R}$ resulting from the QR decomposition is again a triangular matrix, so

$$\tan\theta_{\hat{u},k} \leq \frac{\sigma_{k+1}}{\sigma_k} \tan\theta_{\hat{v},k},$$

provided that $\hat{V}_{11}$ is non-singular. As $\hat{V} = U$,

$$\tan\theta_{\hat{u},k} \leq \frac{\sigma_{k+1}}{\sigma_k} \tan\theta_{u,k}.$$

Hence, the respective (tangents of the) angles between the invariant subspaces and the canonical spaces decrease during a QR decomposition $R^T = Q\hat{R}$ if the corresponding singular values are distinct.

The derivation of the above inequalities for the tangents rests on the relations between left and right singular vector matrices established in the preceding section. These relations, in turn, require that leading principal submatrices of the singular vector matrices be non-singular. It is well-known that the strong non-singularity of the eigenvector matrices is a necessary condition for the convergence to a diagonal matrix of the QR and SVD algorithms [23, 27, 28, 29, 30]. We will now show that a QR decomposition $R^T = Q\hat{R}$ preserves the non-singularity. Although the preservation of strong non-singularity follows from the proofs for the power method, we briefly prove it here from first principles.

If R is non-singular and $R^T = Q\hat{R}$, where Q is orthogonal, then $\hat{R}$ is non-singular. Also, one of the relations between left and right singular vector matrices of R is $U_{11}^T R_{11} = \Sigma_1 V_{11}^T$, so that U_{11} is non-singular whenever V_{11} is. The analogous relation for $\hat{R}$ is $\hat{U}_{11}^T \hat{R}_{11} = \Sigma_1 U_{11}^T$, where $\hat{V} = U$, and the non-singularity of $\hat{R}$ and U_{11} implies the non-singularity of $\hat{V}_{11}$. Thus, if a leading principal submatrix of the original matrix is non-singular, so will be the corresponding leading principal submatrix in each iterate. We summarise the results of this section in the following theorem.

Theorem 3.4.1 *Let R be a non-singular upper triangular matrix of order n with SVD $R = U\Sigma V^T$, singular values*

$$\sigma_1 \geq \ldots \geq \sigma_k \geq \sigma_{k+1} \geq \ldots \geq \sigma_n,$$

and non-singular leading principal submatrices of order k in U and V.

If the upper triangular matrix $\hat{R}$ from the QR decomposition $R^T = Q\hat{R}$ has the SVD $\hat{R} = \hat{U}\Sigma\hat{V}^T$ then

$$\tan\theta_{\hat{u},k} \leq \frac{\sigma_{k+1}}{\sigma_k} \tan\theta_{u,k}, \qquad \tan\theta_{\hat{v},k} \leq \frac{\sigma_{k+1}}{\sigma_k} \tan\theta_{v,k},$$

where $\theta_{z,k}$ is the angle between the leading k columns of a matrix Z and the canonical subspace $\begin{pmatrix} I_k \\ 0 \end{pmatrix}$.

The leading principal submatrices of order k in $\hat{U}$ and $\hat{V}$ are also non-singular.

3.5 Convergence of Singular Vectors

The monotonic convergence of the singular vector matrices and the rates of convergence follow immediately from the results of the previous sections.

Let $R^{(0)}$ be a non-singular upper triangular matrix with SVD $R^{(0)} = U^{(0)}\Sigma V^{(0)^T}$. If, for some k, the leading principal submatrices of order k of $V^{(0)}$ and $U^{(0)}$ are non-singular then

$$\tan\theta_{v,k}^{(i+1)} \leq \frac{\sigma_{k+1}}{\sigma_k}\tan\theta_{v,k}^{(i)}, \qquad \tan\theta_{u,k}^{(i+1)} \leq \frac{\sigma_{k+1}}{\sigma_k}\tan\theta_{u,k}^{(i)},$$

where $\theta_{v,k}^{(i)}$ is the angle between the canonical subspace $\begin{pmatrix} I_k \\ 0 \end{pmatrix}$ and the space spanned by the leading k columns of the singular vector matrix $V^{(i)}$ in the ith iteration; and $\theta_{u,k}^{(i)}$ is the analogous angle for $U^{(i)}$. If, in addition, $\sigma_{k+1} < \sigma_k$ then the leading k columns of the singular vector matrices $V^{(i)}$ and $U^{(i)}$ converge monotonically to canonical form in the above sense.

We can now draw the following conclusions. If all singular values of the matrix $R^{0)}$ are distinct and if the singular vector matrices are strongly non-singular, then the singular vector matrices converge to the identity matrix, in the above sense, monotonically at the rate $\min_k \sigma_{k+1}/\sigma_k$. Since the limit of the singular vector matrices is the identity matrix, the singular values appear in sorted order along the diagonal of the iterates $R^{(i)}$, which converge to a diagonal matrix.

As for the case of multiple singular values, suppose that there is a singular value of multiplicity m, say

$$\sigma_1 \geq \ldots \sigma_k > \sigma_{k+1} = \ldots = \sigma_{k+m} > \sigma_{k+m+1} \geq \ldots \geq \sigma_n.$$

According to the above results for distinct singular values, the singular vector matrices converge to the canonical form

$$\begin{array}{c} \\ k \\ \\ \\ \\ \\ \end{array}\overset{\displaystyle k}{\begin{pmatrix} X & X & & & & \\ X & X & & & & \\ & & X & X & X & X \\ & & X & X & X & X \\ & & X & X & X & X \\ & & X & X & X & X \end{pmatrix}}$$

at the rate σ_{k+1}/σ_k and to the canonical form

$$\begin{array}{c} k \\ + \\ m \\ \\ \\ \\ \end{array}\overset{\displaystyle k \; + \; m}{\begin{pmatrix} X & X & X & X & & \\ X & X & X & X & & \\ X & X & X & X & & \\ X & X & X & X & & \\ & & & & X & X \\ & & & & X & X \end{pmatrix}}$$

at the rate $\sigma_{k+m+1}/\sigma_{k+m}$ (the 'X' represent the non-zero structure of the matrix). Thus the singular vector matrices converge to the canonical form

$$\begin{array}{cc} & \begin{array}{cccccc} k & & m & & & \end{array} \\ \begin{array}{c} k \\ \\ m \\ \\ \\ \\ \end{array} & \begin{pmatrix} X & X & & & & \\ X & X & & & & \\ & & X & X & & \\ & & X & X & & \\ & & & & X & X \\ & & & & X & X \end{pmatrix} \end{array}$$

at the rate $\max\{\sigma_{k+1}/\sigma_k,\ \sigma_{k+m+1}/\sigma_{k+m}\}$.

In general, the singular vector matrices converge to a block-diagonal matrix whose diagonal blocks are orthogonal. The convergence rate is equal to the largest ratio of adjacent distinct singular values. The size of the kth diagonal block equals the multiplicity of the kth distinct singular value, and the columns making up the block represent an orthogonal basis for the associated invariant subspace.

In the limit we can partition the iterate

$$R^{(\infty)} = \begin{pmatrix} R_{11}^{(\infty)} & X & X \\ & R_{22}^{(\infty)} & X \\ & & R_{33}^{(\infty)} \end{pmatrix}$$

into blocks whose sizes conform to the multiplicities, i.e. $R_{11}^{(\infty)}$ is of order k and $R_{22}^{(\infty)}$ is of order m. The singular vector decomposition is partitioned in the same way,

$$U^{(\infty)} = \begin{pmatrix} U_{11}^{(\infty)} & & \\ & U_{22}^{(\infty)} & \\ & & U_{33}^{(\infty)} \end{pmatrix}, \qquad V^{(\infty)} = \begin{pmatrix} V_{11}^{(\infty)} & & \\ & V_{22}^{(\infty)} & \\ & & V_{33}^{(\infty)} \end{pmatrix},$$

where the diagonal blocks U_{ii} and V_{ii} are orthogonal and

$$\Sigma = \begin{pmatrix} \Sigma_1 & & \\ & \sigma_{k+1} I_m & \\ & & \Sigma_{33} \end{pmatrix}, \qquad \Sigma_1 = \begin{pmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_k \end{pmatrix}, \qquad \Sigma_3 = \begin{pmatrix} \sigma_{k+m+1} & & \\ & \ddots & \\ & & \sigma_n \end{pmatrix}.$$

Then $R_{22}^{(\infty)} V_{22}^{(\infty)} = \sigma_{k+1} U_{22}^{(\infty)}$, so that $R_{22}^{(\infty)} = \sigma_{k+1} U_{22}^{(\infty)} {V_{22}^{(\infty)}}^T$ is a multiple of an orthogonal matrix. But $R_{22}^{(\infty)}$ is also upper triangular. Therefore $R_{22}^{(\infty)} = \sigma_{k+1} I_m$ is a scalar matrix, and $U_{22} = V_{22}$ (where we have assumed that $R^{(\infty)}$ has positive diagonal elements).

Therefore, once the columns of each singular vector matrices have converged to a basis for an invariant subspace associated with a multiple singular value, the singular values appear on the diagonal of the iterate $R^{(\infty)}$.

3.6 Convergence Rate for Singular Values

From the convergence rate of the singular vector matrices we can in turn estimate the convergence rate for the singular values.

Suppose again that V_{11} of order k is non-singular. The SVD $R = U\Sigma V^T$ gives

$$R_{22} = U_{21}\Sigma_1 V_{21}^T + U_{22}\Sigma_2 V_{22} = U_{22}(\Sigma_2 + U_{22}^{-1} U_{21}\Sigma_1 V_{21} V_{22}^{-T}) V_{22}^T.$$

Hence, with the results from Section 3.3,

$$||R_{22}|| \leq \sigma_{k+1} + \sigma_1 \tan\theta_{u,k} \tan\theta_{v,k}.$$

Using

$$\tan\theta_{u,k} \leq \frac{\sigma_{k+1}}{\sigma_k} \tan\theta_{v,k}$$

yields

$$\frac{||R_{22}|| - \sigma_{k+1}}{\sigma_{k+1}} \leq \frac{\sigma_1}{\sigma_k} \tan^2\theta_{v,k}.$$

This implies that the relative distance of $||R_{22}||$ from σ_{k+1} is bounded above by the condition number of Σ_1 and the square of the angle between right singular vectors and a canonical subspace. Hence if V_{11} is well-conditioned and the spread of singular values in Σ_1 is small then $||R_{22}||$ is close to σ_{k+1}.

In the same way we derive from $R^{-1} = V\Sigma^{-1}U^T$ that

$$\frac{||R_{11}^{-1}|| - \frac{1}{\sigma_k}}{\frac{1}{\sigma_k}} \leq \frac{\sigma_{k+1}}{\sigma_n} \tan^2\theta_{v,k},$$

and the relative distance of $||R_{11}^{-1}||$ from $1/\sigma_k$ is bounded above by the condition number of Σ_2 and the square of the angle between right singular vectors and a canonical subspace.

In order to determine the rate of convergence, let $R^T = Q\hat{R}$. Since

$$\frac{||\hat{R}_{22}|| - \sigma_{k+1}}{||R_{22}|| - \sigma_{k+1}} \approx \left(\frac{\tan\theta_{\hat{v},k}}{\tan\theta_{v,k}}\right)^2 \leq \left(\frac{\sigma_{k+1}}{\sigma_k}\right)^2$$

and

$$\frac{||\hat{R}_{11}^{-1}|| - \frac{1}{\sigma_k}}{||R_{11}^{-1}|| - \frac{1}{\sigma_k}} \approx \left(\frac{\tan\theta_{\hat{v},k}}{\tan\theta_{v,k}}\right)^2 \leq \left(\frac{\sigma_{k+1}}{\sigma_k}\right)^2,$$

the rate of convergence of principal submatrices with disjoint singular values is the square of that of the associated singular vectors.

In the case of a singular values σ_{k+1} of multiplicity $m > 1$,

$$\sigma_1 \geq \ldots \sigma_k > \sigma_{k+1} = \ldots = \sigma_{k+m} > \sigma_{k+m+1} \geq \cdots \geq \sigma_n,$$

partition the iterates as in the previous section,

$$R^{(i)} = \begin{pmatrix} R_{11}^{(i)} & X & X \\ & R_{22}^{(i)} & X \\ & & R_{33}^{(i)} \end{pmatrix}.$$

The convergence of $||R_{11}^{(i)^{-1}}||$ to $1/\sigma_k$ and of

$$\left\| \begin{pmatrix} R_{22}^{(i)} & X \\ & R_{33}^{(i)} \end{pmatrix} \right\|$$

to σ_{k+1} occurs at the rate $\sigma_{k+1}^2/\sigma_k^2$, while the convergence of

$$\left\| \begin{pmatrix} R_{11}^{(i)} & X \\ & R_{22}^{(i)} \end{pmatrix}^{-1} \right\|$$

to $1/\sigma_{k+1}$ and the convergence of $R_{33}^{(i)}$ to σ_{k+m+1} occurs at the rate $\sigma_{k+1}^2/\sigma_{k+m+1}^2$.

Hence a principal submatrix $R_{22}^{(i)}$ associated with a singular value σ_{k+1} of multiplicity m converges to $\sigma_{k+1}I_m$ at the rate $\max\{\sigma_{k+1}^2/\sigma_k^2, \sigma_{k+m+1}^2/\sigma_{k+1}^2\}$.

3.7 Summary

We summarise all convergence results for the algorithm

$$R^{(i)^T} = Q^{(i+1)} R^{(i+1)}, \qquad i \geq 0,$$

in the following theorem.

Theorem 3.7.1 *Let $R^{(0)}$ be an upper triangular matrix of order n with SVD $R = U^{(0)} \Sigma V^{(0)^T}$ and singular values*

$$\sigma_1 \geq \ldots \sigma_k > \sigma_{k+1} = \ldots = \sigma_{k+m} > \sigma_{k+m+1} \geq \cdots \geq \sigma_n;$$

and let $R^{(i)} = U^{(i)} \Sigma V^{(i)^T}$ be the SVDs of the iterates.

If $R^{(0)}$ is non-singular, and if $U^{(0)}$ and $V^{(0)}$ are strongly non-singular then the following convergence results hold.

$m = 1$: *Columns $k+1$ of $U^{(i)}$ and $V(i)$ converge to column $k+1$ of the identity matrix at the rate*

$$\rho_k = \max\{\frac{\sigma_{k+1}}{\sigma_k}, \frac{\sigma_{k+2}}{\sigma_{k+1}}\},$$

and the $(k+1)$st diagonal element of $R^{(i)}$ converges to σ_{k+1} at the rate ρ_k^2.
Convergence of the singular vectors is monotonic in the sense that

$$\tan \theta_{v,k}^{(i+1)} \leq \frac{\sigma_{k+1}}{\sigma_k} \tan \theta_{v,k}^{(i)}, \qquad \tan \theta_{u,k}^{(i+1)} \leq \frac{\sigma_{k+1}}{\sigma_k} \tan \theta_{u,k}^{(i)},$$

where $\theta_{v,k}^{(i)}$ is the angle between the canonical subspace $\begin{pmatrix} I_k \\ 0 \end{pmatrix}$ and the space spanned by the leading k columns of $V^{(i)}$; and $\theta_{u,k}^{(i)}$ is the analogous angle for $U^{(i)}$.
Convergence of the singular values is monotonic in the sense that

$$\|R_{11}^{(i+1)^{-1}}\| \leq \|R_{11}^{(i)^{-1}}\|, \qquad \|R_{22}^{(i+1)}\| \leq \|R_{22}^{(i)}\|$$

and

$$\frac{\|R_{22}^{(i)}\| - \sigma_{k+1}}{\sigma_{k+1}} \leq \frac{\sigma_1}{\sigma_k} \tan^2 \theta_{v,k}^{(i)}, \qquad \frac{\|R_{11}^{(i)^{-1}}\| - \frac{1}{\sigma_k}}{\frac{1}{\sigma_k}} \leq \frac{\sigma_{k+1}}{\sigma_n} \tan^2 \theta_{v,k}^{(i)}.$$

$m > 1$: *Columns $k+1, \ldots, k+m$ of $U^{(i)}$ and $V^{(i)}$ converge to a $n \times m$ matrix of the form*

$$\begin{array}{c} \\ k \\ m \\ \\ \end{array} \begin{array}{c} m \\ \begin{pmatrix} 0 \\ Z \\ 0 \end{pmatrix} \end{array},$$

where Z is orthogonal, at the rate

$$\rho_k = \max\{\frac{\sigma_{k+1}}{\sigma_k}, \frac{\sigma_{k+m+1}}{\sigma_{k+1}}\}.$$

The principal submatrix of order m of $R^{(i)}$,

$$\begin{pmatrix} R_{k+1,k+1}^{(i)} & \cdots & R_{k,k+m}^{(i)} \\ & \ddots & \vdots \\ & & R_{k+m,k+m}^{(i)} \end{pmatrix},$$

converges to $\sigma_{k+1} I_m$ at the rate ρ_k^2.

3.8 Consequences

The upper bounds on the relative distance between $\|R_{11}^{(i)^{-1}}\|$ and $\|R_{22}^{(i)}\|$ to the respective singular values depend on the spreads $\sigma_1 \ldots \sigma_k$ and $\sigma_{k+1} \ldots \sigma_n$, and the conditioning of the leading principal submatrices of order k of $V^{(i)}$. The number of iterations required to reduce the relative distance between $\|R_{22}\|$ and σ_{k+1} to ϵ can thus be estimated as

$$\frac{\log \sigma_1/\sigma_k - \log \epsilon + \log \tan \theta_{v,k}^{(0)}}{\log \sigma_k/\sigma_{k+1}}.$$

An analogous estimate can be made for $\|R_{11}^{-1}\|$.

It is observed in Section 2.2 of [23] that the QR algorithm tends to converge to the small eigenvalues first. According to our analysis, though, there is no preference of the unshifted Golub-Kahan algorithm for small singular values over larger ones. However, such a preference may be enforced by a suitable choice of shifts [22, 30].

The same upper bounds also explain why both algorithms have such a hard time with graded matrices whose elements increase in size towards the bottom, cf. Section 5 in [11] and Theorem 5 in [12]. These matrices have a large spread in the spectrum and very ill-conditioned leading principal submatrices. One of the simplest example of a graded matrix is

$$R_0 = \begin{pmatrix} 1 & \epsilon \\ & \alpha \end{pmatrix},$$

where $\epsilon \ll 1 \ll \alpha$. One iteration of the extended Golub-Kahan algorithm gives

$$R_1 = \frac{1}{\sqrt{1+\epsilon^2}} \begin{pmatrix} 1+\epsilon^2 & \alpha\epsilon \\ & \alpha \end{pmatrix},$$

whose off-diagonal element has increased from ϵ to $\alpha\epsilon$. But the diagonal elements have only changed marginally, and it is obvious that many iterations are needed to arrive at a diagonal matrix with diagonal elements in descending order (a similar example was used in Section 8.7 of [30] to illustrate slow convergence of the LR algorithm). Section 4.2 illustrates how to force fast convergence on such graded matrices without the need to decide between QR and QL-type algorithms as in [11, 12].

Numerically, the leading principal submatrices of the singular vector matrices are usually nonsingular but may be very ill-conditioned, hence the slow convergence. Nonsingularity of the principal submatrices happens due to finite precision arithmetic, which causes small perturbations of zero off-diagonal elements, as in the example above.

4 Rank-Revealing QR and URV Decompositions

In this section we discuss the connections between our algorithm, and rank-revealing QR (RRQR) decompositions [7] and the URV decomposition [25].

4.1 Preprocessing by Rank-Revealing QR Decompositions

As before, let

$$R = \begin{array}{c} \\ k \end{array}\!\!\begin{array}{c} k \\ \begin{pmatrix} R_{11} & R_{12} \\ & R_{22} \end{pmatrix} \end{array}$$

be a non-singular upper triangular matrix. For each k define

$$\gamma_k = \|R_{11}^{-1}\|\,\|R_{22}\|.$$

When $\gamma_k < 1$ then

$$\|R_{22}\| < \frac{1}{\|R_{11}^{-1}\|},$$

which means that all singular values of R_{11} are larger than the singular values of R_{22}, and a partial ordering of the singular values of R has occurred: the k largest singular values of R are represented by R_{11} and separated from the remaining smaller ones, which are represented by R_{22}. That is why we refer to γ_k as the *separation* between R_{11} and R_{22}.

Given a specific k (usually determined by the number of singular values of R that are smaller than a certain threshold), it is the objective of RRQR algorithms to find a permutation matrix P so that in the QR decomposition $RP = Q\bar{R}$ the resulting triangular matrix $\bar{R}$ has a $(1,1)$ block with maximal smallest singular value, and/or a (2,2) block with a minimal largest singular value [7]. That is, the goal of RRQR algorithms is to make the separation between the singular values of R_{11} and those of R_{22} as large as possible, thus to minimise γ_k for a particular k.

Although the tangent of the angle between singular vectors and a canonical subspace converges monotonically, this cannot in practice be monitored cheaply during the course of our algorithm. Therefore we express the convergence behaviour in terms of observable quantities, namely the separation γ_k. According to Section 3.2, γ_k decreases monotonically for each k,

$$\gamma_k^{(i+1)} \leq \gamma_k^{(i)},$$

so the separation between singular values of $R_{11}^{(i)}$ and $R_{22}^{(i)}$ never decreases. Section 3.6 showed that if the leading principal submatrix of the singular vector matrices is non-singular and if $\sigma_{k+1}/\sigma_k < 1$ then

$$\|R_{11}^{(i)^{-1}}\| \rightarrow 1/\sigma_k, \qquad \|R_{22}^{(i)}\| \rightarrow \sigma_{k+1}, \qquad \gamma_k^{(i)} \rightarrow \sigma_{k+1}/\sigma_k \quad \text{as} \quad i \rightarrow \infty.$$

Because the convergence of $\gamma_k^{(i)}$ to $\sigma_{k+1}/\sigma_k < 1$ is monotone, there exists a number i_k such that $\gamma_k^{(i)} < 1$ for all $i \geq i_k$. It make sense therefore to distinguish, for each k, two phases of our algorithm depending on the value of $\gamma_k^{(i)}$:

1. a rank-revealing phase, where $\gamma_k^{(i)} > 1$, during which the singular values of $R_{11}^{(i)}$ and $R_{22}^{(i)}$ separate; and
2. a monotonic phase, where $\gamma_k^{(i)} \leq 1$, during which *all* quantities of interest converge monotonically.

We will show that once the monotonic phase has been reached for some k, the iterates $R^{(i)}$ converge rapidly to block-diagonal form because the off-diagonal blocks decrease monotonically.

4.2 The Rank-Revealing Phase

Theoretically one can enforce the onset of the monotonic phase for a particular k in a finite number of operations, provided the singular values σ_k and σ_{k+1} are well-separated. This is done by preceding our algorithm with a good RRQR algorithm, which implements the rank-revealing phase. Moreover, a preliminary RRQR algorithm can also reverse the grading in a matrix all of whose large elements are at the bottom of the matrix. This obviates the need for deciding whether to subject the matrix to an algorithm of QR or QL type [12, 11].

The idea of permuting rows or columns of the iterates during eigenvalue computations is not new. Pivoting, in the form of row exchanges, has been suggested for the LR algorithm, Section 8.13 in [30] and Section 2.7 in [23], to enhance numerical stability in those cases where the orthodox LR algorithm fails to converge.

The most accurate RRQR algorithm known so far is Hybrid III(k) from [7]. Given an index k and a matrix R of order n with singular values

$$\sigma_1 \geq \ldots \geq \sigma_k \geq \sigma_{k+1} \geq \ldots \geq \sigma_n,$$

Hybrid III(k) finds a permutation matrix P so that the blocks of the triangular matrix

$$\bar{R} = \begin{matrix} & k & \\ k & \begin{pmatrix} \bar{R}_{11} & \bar{R}_{12} \\ & \bar{R}_{22} \end{pmatrix} \end{matrix}$$

in $RP = Q\bar{R}$ satisfy

$$\frac{1}{\|\bar{R}_{11}^{-1}\|} \geq \frac{1}{\sqrt{k(n-k+1)}}\sigma_k, \qquad \|\bar{R}_{22}\| \leq \sqrt{(k+1)(n-k)}\sigma_{k+1}.$$

Hybrid-III(k) guarantees bounds for both, the (1,1) and the (2,2) block, so that

$$\gamma_k = \|\bar{R}_{11}^{-1}\|\,\|\bar{R}_{22}\| \leq (k+1)(n-k+1)\,\frac{\sigma_{k+1}}{\sigma_k}.$$

In practice, it would probably suffice to precede our algorithm with a cheaper and possibly less accurate form of column pivoting (an attempt at explaining the practical effectiveness of the simple column pivoting strategies, regardless of their potential failures, is made in [7]). We briefly discuss the three types of simple pivoting strategies on which the existing RRQR algorithms are based.

The algorithm by Golub [4, 14, 18], which is known as *the* QR decomposition with column pivoting; algorithm Greedy-I.4 in [7]; and the one by Chan and Hansen [6] all assemble the columns with largest norm in the left part of the matrix. They guarantee only exponential bounds on $\|\bar{R}_{11}^{-1}\|$,

$$\frac{1}{\|\bar{R}_{11}^{-1}\|} \geq \frac{\sigma_k}{n2^k}, \qquad \gamma_k \leq n\,2^k\,\frac{\|\bar{R}_{22}\|}{\sigma_k}.$$

The algorithms by Gragg and Stewart [19], Stewart [24], Chan [5], and Foster [13] assemble the rows of the inverse with smallest norm in the lower part of the matrix. They guarantee only exponential bounds on $\|\bar{R}_{22}\|$,

$$\|\bar{R}_{22}\| \leq n\,2^{n-k}\sigma_{k+1}, \qquad \gamma_k \leq n\,2^{n-k}\,\sigma_{k+1}\,\|\bar{R}_{11}^{-1}\|.$$

The algorithm by Golub, Klema and Stewart [16] permutes the rows of the right singular vector matrix

$$V = \begin{matrix} & k & \\ k & \begin{pmatrix} V_{11} & V_{12} \\ V_{21} & V_{22} \end{pmatrix} \end{matrix}$$

so that V_{11} becomes well-conditioned; this is accomplished by applying the QR decomposition with column pivoting [4, 14, 18] to $(V_{11}^T \quad V_{21}^T)$. The algorithm guarantees exponentiell bounds on both $\|\bar{R}_{11}^{-1}\|$ and $\|\bar{R}_{22}\|$,

$$\frac{1}{\|\bar{R}_{11}^{-1}\|} \geq \frac{\sigma_k}{n 2^{\min\{k,n-k\}}}, \qquad \|\bar{R}_{22}\| \leq n\, 2^{\min\{k,n-k\}} \sigma_{k+1}, \qquad \gamma_k \leq n^2\, 2^{2\min\{k,n-k\}} \frac{\sigma_{k+1}}{\sigma_k}.$$

4.3 A Divide-and-Conquer Algorithm

We will show in Section 4.4 that once the rank-revealing phase has been completed for some k, the iterates converge rapidly to diagonal form because the off-diagonal blocks are decreasing monotonically and deflation can be expected to occur fast. Hence, preceding our algorithm with a RRQR algorithm for some k corresponds to forcing completion of the rank-revealing phase and the start of deflation for that k.

This observation can be exploited to design a divide-and-conquer algorithm to compute singular values of dense matrices A. Such an algorithm may be advantageous on a parallel architecture. Below is a rough sketch.

1. Select a k and perform a RRQR algorithm $AP = Q\bar{R}$ so that $\gamma_k = \|\bar{R}_{11}^{-1}\|\,\|\bar{R}_{22}\| < 1$.
2. Apply our algorithm to $R^{(0)} = \bar{R}$ and iterate until $\|R_{12}^{(i)}\|$ is small enough.
3. Apply Steps 1 and 2 recursively to $R_{11}^{(i)}$ and to $R_{22}^{(i)}$.

The index k selected in step 1 determines where the matrix is to be divided. There are several ways to determine k. The simplest option is to set $k = n/2$ and choose Hybrid III(n/2) as the RRQR algorithm. This would break the matrix into equally sized blocks and ensure load balance with regard to parallel execution. But the separation of the singular values $\sigma_{n/2}$ and $\sigma_{n/2+1}$ may not be large enough.

To circumvent this problem one could alternatively run Golub's QR with column pivoting on the matrix and select as k that index j for which the estimate $|\bar{r}_{jj}|/|\bar{r}_{j+1,j+1}|$ of γ_j is smallest. Instead of using $\bar{r}_{jj}$ one could also estimate the norm of $\|\bar{R}_{jj}^{-1}\|$ by an incremental condition estimator [1, 2]. We do not have enough computational experience to judge whether the algorithm presents a viable alternative to other methods that operate on dense matrices, such as a Jacobi-type method, e.g. [3, 8].

4.4 The Monotonic Phase

Now we will show that once the monotonic phase has been reached, convergence to block-diagonal form is very fast because the off-diagonal blocks $R_{12}^{(i)}$ of the iterates decrease monotonically once $\gamma_k^{(i)} < 1$.

This can be seen as follows. From $R^T = Q\hat{R}$ and $\hat{R} = Q^T R^T$ it follows that

$$\hat{R}_{12} = Q_{21}^T R_{22}^T, \qquad R_{12}^T = Q_{21}\hat{R}_{11}, \qquad R_{11}^T = Q_{11}\hat{R}_{11}.$$

Thus, $\hat{R}_{12} = Q_{11}^T R_{11}^{-1} R_{12} R_{22}^T$ and

$$\|\hat{R}_{12}\| \leq \gamma_k \|R_{12}\|.$$

Thus, during the monotonic phase

$$\|R_{12}^{(i+1)}\| \leq \gamma_k^{(i)} \|R_{12}^{(i)}\| < \|R_{12}^{(i)}\|.$$

Since $\gamma_k^{(i)} \to \sigma_{k+1}/\sigma_k$, the blocks corresponding to well-separated singular values may decrease faster and deflation (i.e. the decoupling of the (1,1) and (2,2) blocks into two separate problems due to an almost zero off-diagonal block) is likely to set in earlier.

4.5 Postprocessing with Our Algorithm

In the previous sections we discussed how to use RRQR algorithms as preprocessors to enforce monotonic convergence to block-diagonal form in our algorithm. Now we look at the other side of the coin to study how the our algorithm can be used as a postprocessor to refine a more general form of RRQR factorisation, called URV factorisation.

The URV factorisation was introduced by Stewart [25] to compute the null spaces of a matrix that is repeatedly updated. If R is a real $n \times n$ matrix of rank k then there exist orthogonal matrices U and V of order n and a non-singular upper triangular matrix $\bar{R}$ of order k such that

$$R = U \begin{pmatrix} \bar{R} & 0 \\ 0 & 0 \end{pmatrix} V^T.$$

The URV decomposition constitutes a partial SVD, and it can be viewed as a compromise between RRQR factorisation and SVD.

In applications it is often the case that R is almost of rank k, that is, its singular values σ_{k+1}, ..., σ_n are small. In this case one would like to find a decomposition

$$R = U \begin{pmatrix} \bar{R}_{11} & \bar{R}_{12} \\ & \bar{R}_{22} \end{pmatrix} V^T,$$

where $\|\bar{R}_{11}^{-1}\|$ is maximal, i.e. close to $1/\sigma_k$; where $\|\bar{R}_{22}\|$ is minimal, i.e. close to σ_{k+1}; and where $\bar{R}_{12}$ is small, i.e. close to σ_{k+1}. Hence, the URV decomposition differs from a RRQR factorisation by the requirement that $\|\bar{R}_{12}\|$ also be small. For this purpose, rotations rather than just permutations from the right are allowed.

In [25, 26] Stewart proposes to compute the URV decomposition as follows: compute orthogonal matrices P and Q such that $RP = Q\bar{R}$ where $\|\bar{R}_{11}^{-1}\|$ is close to $1/\sigma_k$ and $\|(\bar{R}_{12}^T \quad \bar{R}_{22})\|$ is close to σ_{k+1}. Then perform several of the following 'refinement steps' on $R^{(0)} = \bar{R}$ to further decrease the size of the (1,2) block [26]: first determine an orthogonal matrix $Q^{(1)}$ so that $R^{(1)^T} = R^{(0)}Q^{(1)}$ is lower triangular and second determine an orthogonal matrix $Q^{(2)}$ so that $R^{(2)} = Q^{(2)^T} R^{(1)^T}$ is upper triangular. In [25] Stewart proposes an incomplete version of these refinement steps: reduce only the last column of $R(0)$ to e_n, and in this resulting matrix in turn reduce only the last row to e_n^T.

Note that in the beginning these algorithms accomplish more than a RRQR decomposition. Due to the rotations performed on both sides of the matrix the off-diagonal block will also be small. Hence the following result from [26] applies: if the off-diagonal block is small enough, i.e. if

$$\|R_{12}^{(0)}\| + \|R_{22}^{(0)}\| < \sigma_k$$

then the first part of the refinement steps in [25, 26] causes a monotonic decrease $\|R_{12}^{(1)}\| < \|R_{12}^{(0)}\|$ in the (1,2) block, and so does, of course, the second part of the refinement step.

However, the refinement step in [26] is nothing else but two iterations of our algorithm

$$R^{(0)^T} = Q^{(1)}R^{(1)}, \qquad R^{(1)^T} = Q^{(2)}R^{(2)}.$$

The refinement step in [25] amounts to performing incomplete QR factorisations, where $R_{11}^{(1)}$ remains lower triangular.

As we showed in the previous section, no assumption on the (1,2) block is needed in general to ensure monotonic decrease, as long the singular values of $R_{11}^{(i)}$ and $R_{22}^{(i)}$ have been separated: if $\gamma_k^{(i)} < 1$ then $\|R_{12}^{(i+1)}\| < \|R_{12}^{(i)}\|$. This is true regardless of whether σ_{k+1} is small or not. However, if $\|R_{22}^{(i)}\|$ is small then $\|R_{12}^{(i+1)}\|$ is as small, – regardless of the relation between $R_{11}^{(i)}$ and $R_{22}^{(i)}$ – because $R_{12}^{(i+1)} = Q_{21}^{(i+1)}R_{22}^{(i)^T}$, so $\|R_{12}^{(i+1)}\| \leq \|R_{22}^{(i)}\|$. Hence, one can compute a URV decomposition of R by determining a RRQR decomposition $RP = Q\bar{R}$ and then performing several steps of our algorithm on $\bar{R}$, which will converge monotonically to the desired URV decomposition.

5 Deflation and Convergence Criteria

Demmel and Kahan [11] have shown that, in finite precision arithmetic, the Golub-Kahan algorithm for bi-diagonal matrices computes singular values to high relative accuracy, provided a zero shift is used for the smallest singular values and the algorithm is implemented without subtractions. Fernando and Parlett [12] modify Rutishauser's differential QD algorithm for bi-diagonal matrices to obtain an algorithm that is faster than Demmel's and Kahan's implementation of the Golub-Kahan algorithm but still computes singular values to high relative accuracy.

5.1 Bi-Diagonal Matrices

Demmel and Kahan provide two types of convergence criteria that preserve high relative accuracy of the computed singular values.

Convergence Criterion 1a sets the off-diagonal block R_{12} in R to zero whenever $\|R_{12}\|\,|e^T R_{11}^{-1} e_k|$ is small enough, where e is the vector of all ones and e_k is the kth column of the identity matrix. Note that $\|R_{12}\|$ is just the absolute value of the off-diagonal element in an upper bi-diagonal matrix. If R has also positive diagonal elements and negative off-diagonal elements, which can always be accomplished by multiplying R with orthogonal diagonal matrices, then $|e^T R_{11}^{-1} e_k|$ represents a lower bound for the one-norm of R_{11}^{-1},

$$|e^T R_{11}^{-1} e_k| \leq \|e^T R_{11}^{-1}\|_1 = \|R_{11}^{-1}\|_1 \leq \sqrt{k}\|R_{11}^{-1}\|.$$

Thus, Convergence Criterion 1a has the upper bound

$$\|R_{12}\|\,|e^T R_{11}^{-1} e_k| \leq \sqrt{k}\,\|R_{11}^{-1}\|\,\|R_{12}\|.$$

Convergence Criterion 1b is the equivalent of Convergence Criterion 1a applied to an algorithm based on a QL decomposition, i.e.a decomposition into a product of orthogonal and lower triangular matrix.

Convergence Criterion 2a sets the off-diagonal block to zero whenever

$$\frac{2\|R_{12}\|^2}{\text{gap}_k(\sigma_{min}(R_{11}) + \|R_{22}\|)}$$

is small enough, where $\sigma_{min}(R_{11}) = 1/\|R_{11}^{-1}\|$ and $\text{gap}_k = \sigma_{min}(R_{11}) - \|R_{22}\|$. Convergence Criterion 2b is the analogue for a QL-based algorithm.

Demmel and Kahan show that the application of these convergence criteria causes essentially only a relative perturbation in the singular values. Suppose the singular values of R_{11} and R_{22} are written as $\mu_1 \geq \ldots \geq \mu_n$, and let $0 < \eta < 1$ be the relative accuracy to which the singular values are to be computed.

As for Convergence Criteria 1a/b, Theorem 4 in [11], if $\|R_{12}\|\,|e^T R_{11}^{-1} e_k| < \eta$ then each singular value σ_j of the bi-diagonal matrix R satisfies

$$-m\phi(\eta) \leq \ln \frac{\mu_j}{\sigma_j} \leq m\phi(\eta),$$

where $\phi(\eta) \leq \eta/\sqrt{2}$, and there are at most m singular values μ_l, one of them being μ_j, whose intervals

$$\left\{\mu : \quad -\phi(\eta) \leq \ln \frac{\mu}{\mu_l} \leq \phi(\eta)\right\}$$

overlap.

As for Convergence Criterion 2a, Theorem 5 in [11], if

$$\frac{2\|R_{12}\|^2}{\text{gap}_k(\sigma_{min}(R_{11}) + \|R_{22}\|)} < \eta$$

then each singular value σ_j of the bi-diagonal matrix R satisfies $|\sigma_j - \mu_j| \approx m\eta\mu_j$, where $m\eta \ll 1$ and there are at most m singular values μ_l, with μ_j among them, whose intervals

$$\{\mu : \quad |\mu_l - \mu| \leq \eta\mu\}$$

overlap.

5.2 Triangular Matrices

We now extend the two types of convergence criteria in [11] to the computation of singular values for triangular matrices.

The extension to triangular matrices is accomplished by generalising the following theorem of Stewart from [25]. If $\|R_{12}\| + \|R_{22}\| \leq \sigma_k$ then

$$\frac{|\sigma_{k+i} - \sigma_i(R_{22})|}{\sigma_1(R_{22})} \leq \frac{\|R_{12}\|\,\|R_{22}\|}{\delta^2 - \|R_{22}\|^2}, \qquad \delta = \sigma_k - \|R_{12}\|.$$

We now derive a simpler bound that holds without any assumptions on the size of $\|R_{12}\|$. To simplify notation, set $R = R^{(0)}$ and $\gamma_k = \gamma^{(0)}$.

From $R^T = Q\hat{R}$ it follows that

$$\begin{pmatrix} 0 \\ R_{22}^T \end{pmatrix} = Q \begin{pmatrix} 0 \\ \hat{R}_{22} \end{pmatrix} + Q \begin{pmatrix} \hat{R}_{12} \\ 0 \end{pmatrix}.$$

Using $|\sigma_j(A+E) - \sigma_j(A)| \leq \|E\|$, Corollary 8.3.2 in [18], with

$$A = Q \begin{pmatrix} 0 \\ \hat{R}_{22} \end{pmatrix}, \qquad A + E = Q \begin{pmatrix} 0 \\ \hat{R}_{22} \end{pmatrix} + Q \begin{pmatrix} \hat{R}_{12} \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ R_{22}^T \end{pmatrix}$$

yields for the singular values of R_{22} and $\hat{R}_{22}$

$$|\sigma_j(R_{22}) - \sigma_j(\hat{R}_{22})| \leq \|\hat{R}_{12}\|.$$

But in Section 4.4 we showed that

$$\|\hat{R}_{12}\| \leq \gamma_k \|R_{12}\|, \qquad \gamma_k = \|R_{11}^{-1}\| \, \|R_{22}\|.$$

Thus

$$|\sigma_j(R_{22}) - \sigma_j(\hat{R}_{22})| \leq \gamma_k \|R_{12}\|.$$

According to Section 4.1, $\gamma_k^{(i)}$ is monotonically decreasing, $\gamma_k^{(i+1)} \leq \gamma_k^{(i)}$, and the difference between two successive iterations is

$$|\sigma_j(R_{22}^{(i+1)}) - \sigma_j(R_{22}^{(i)})| \leq \|R_{12}^{(i+1)}\| \leq \gamma_k^{(i)} \|R_{12}^{(i)}\| \leq \gamma_k^i \|R_{12}\|.$$

As for the difference between iteration $i+2$ and i, we employ the idea in Stewart's proof in [25],

$$|\sigma_j(R_{22}^{(i+2)}) - \sigma_j(R_{22}^{(i)})| \leq |\sigma_j(R_{22}^{(i+2)}) - \sigma_j(R_{22}^{(i+1)})| + |\sigma_j(R_{22}^{(i+1)}) - \sigma_j(R_{22}^{(i)})| \leq \left(\gamma_k^{i+1} + \gamma_k^i\right) \|R_{12}\|.$$

When V_{11} is non-singular then, from Section 3.6, the singular values of $R_{22}^{(i)}$ converge to the singular values $\sigma_{k+1}, \ldots, \sigma_n$ of Σ_2 as $i \to \infty$. Whenever $\gamma_k < 1$ we get in the limit

$$|\sigma_{k+j} - \sigma_j(R_{22})| \leq \|R_{12}\| \sum_{l=1}^{\infty} \gamma_k^l = \|R_{12}\| \frac{\gamma_k}{1-\gamma_k}$$

since $\sum_{l=1}^{\infty} \gamma_k^l = \frac{1}{1-\gamma_k} - 1$. Using $\gamma_k = \|R_{11}^{-1}\| \, \|R_{22}\|$ and $\sigma_{min}(R_{11}) = 1/\|R_{11}^{-1}\|$ yields

$$\frac{|\sigma_{k+j} - \sigma_j(R_{22})|}{\sigma_1(R_{22})} \leq \frac{\|R_{11}^{-1}\|}{1-\gamma_k} \|R_{12}\| = \frac{\|R_{12}\|}{\sigma_{min}(R_{11}) - \|R_{22}\|}.$$

We summarise the results in the following theorem.

Theorem 5.2.1 *Let*

$$R \;=\; \begin{array}{c} \\ k \end{array} \!\!\begin{array}{c} k \\ \begin{pmatrix} R_{11} & R_{12} \\ & R_{22} \end{pmatrix} \end{array}$$

be a non-singular upper triangular matrix of order n with SVD $R = U\Sigma V^T$ and singular values

$$\sigma_1 \geq \ldots \geq \sigma_k \geq \sigma_{k+1} \geq \ldots \geq \sigma_n.$$

If V is strongly non-singular and if $\text{gap}_k = \sigma_{min}(R_{11}) - \|R_{22}\| > 0$ *then the singular values $\sigma_j(R_{22})$ of R_{22} satisfy*

$$\frac{|\sigma_{k+j} - \sigma_j(R_{22})|}{\sigma_1(R_{22})} \leq \frac{\|R_{12}\|}{\text{gap}_k}.$$

5.3 Consequences

The implications of Theorem 5.2.1 discussed in this section are to be interpreted as statements about the nature of a perturbation result rather than as statements about the accuracy of our algorithm in finite precision. Although Demmel and Kahan [11] have presented an implementation of the Golub-Kahan algorithm for bi-diagonal matrices that guarantees high relative accuracy in finite precision, it is not clear that there can be an implementation that guarantees high relative accuracy for triangular matrices in finite precision.

Theorem 5.2.1 is most valuable for the case $k = n - m$, where m is the multiplicity of the smallest singular value σ_n. It represents a convergence criterion that assures high relative accuracy for the computation of σ_n (this presumes that no premature deflation was enforced in the interior of the matrix at some earlier time). Theorem 5.2.1 requires $\mathrm{gap}_k > 0$, which means that all singular values of R_{11} are larger than those of R_{22}. Since $\mathrm{gap}_k > 0$ is equivalent to $\gamma_k < 1$, Theorem 5.2.1 is applicable to our algorithm only once the monotonic phase for k has set in. Note that $\mathrm{gap}_k > 0$ is not satisfied for a graded matrix whose elements increase in size towards the bottom, regardless of how small R_{12} is.

Theorem 5.2.1 represents the 'square-root of' Convergence Criterion 2a in [11]. It also equals, up to a factor $\sqrt{k}$, the upper bound on Convergence Criterion 1a, which was derived in the previous section. This can be seen from the Taylor series expansion,

$$\frac{\gamma_k}{1-\gamma_k} = \gamma_k + O(\gamma_k^2),$$

so

$$\frac{|\sigma_{k+j} - \sigma_j(R_{22})|}{\sigma_1(R_{22})} \leq \|R_{11}^{-1}\|\,\|R_{12}\| \quad + \|R_{12}\|\,\|R_{22}\|\, O(\gamma_k^2).$$

Note, however, that Convergence Criteria 1a/b in [11] apply regardless of whether $\gamma_k < 1$.

The bound in Theorem 5.2.1 suggests using the simple deflation criterion

$$\|R_{12}\| \leq \eta\, \frac{\mathrm{gap}_k}{\|R_{22}\|}$$

in order to guarantee the computation of singular values of triangular matrices to absolute accuracy η – assuming that no deflation in the interior of the matrix has been enforced earlier. If $\|R_{22}\|$ is small and the singular values of R_{11} and R_{22} are well-separated then this criterion recognises convergence earlier than the traditional criterion, Corollary 8.3.2 in [18],

$$\|R_{12}\| \leq \eta.$$

Relative accuracy η for *all* singular values is achieved if

$$\|R_{12}\| \leq \eta \frac{\mathrm{gap}_k}{\kappa(R_{22})},$$

where $\kappa(R_{22}) = \|R_{22}\|\,\|R_{22}^{-1}\|$ is the condition number of R_{22}.

6 Summary

We have presented a new algorithm for computing the singular value decomposition of real triangular matrices R that avoids the formation of $R^T R$ or RR^T and instead performs repeated QR factorisations on iterates of R.

We have demonstrated the monotonic convergence of the singular vectors, and we showed that the rate of convergence of the singular values is the square of that of the singular vectors. The most important ingredient in these proofs is the exploitation of the structure of the SVD of triangular matrices. The convergence results also explain the occurrence of deflation and the slow convergence for inappropriately graded matrices. These results do not only apply to our algorithm but also to the unshifted QR algorithm for computing eigenvalues. Our analysis can be easily extended to multiple-step algorithms with stationary shifts [29] for the computation of eigenvalues and singular values.

By making the connection to RRQR and URV decompositions we have developed new ideas for several algorithms:

- A preliminary RRQR decomposition preceding the application of our algorithm so as to reverse the grading in inappropriately graded matrices and enhance subsequent convergence.
- The computation of the URV decomposition by several iterations of our algorithm.
- Alternating application of RRQR decompositions and several iterations of our algorithm in order to enforce premature deflation of off-diagonal blocks in a divide-and-conquer algorithm for computing singular values of dense matrices.

We have presented a very simple convergence and deflation criterion for triangular matrices that recognises convergence of the singular values earlier than the traditional perturbation bounds. In particular, it permits high relative accuracy in the smallest singular values. This convergence criterion represents a generalisation of some convergence criteria for bi-diagonal matrices in [11]. A more careful perturbation analysis may lead to more optimistic results about deflation in the interior of the matrix.

Since we have only presented ideas in this paper, the next step will, of course, be to gather numerical evidence to determine whether our ideas give rise to numerically viable algorithms.

Acknowledgements

We would like to thank Stan Eisenstat, Beresford Parlett, and David Watkins for helpful discussions.

References

[1] C. Bischof, *Incremental condition estimation*, SIAM J. Matrix Anal. Appl., 11 (1990), pp. 312–22.

[2] C. Bischof, J. Lewis, and D. Pierce, *Incremental condition estimation for sparse matrices*, SIAM J. Matrix Anal. Appl., 11 (1990), pp. 644–59.

[3] R. Brent and F. Luk, *The solution of singular-value and symmetric eigenvalue problems on multiprocessor arrays*, Siam J. Sci. Stat. Comput., 6 (1985), pp. 69–84.

[4] P. Businger and G. Golub, *Linear least squares solutions by Householder transformations*, Numer. Math., 7 (1965), pp. 269–76.

[5] T. Chan, *Rank revealing QR factorizations*, Linear Algebra and its Applications, 88/89 (1987), pp. 67–82.

[6] T. Chan and P. Hansen, *Low-rank revealing QR factorizations*, Tech. Report 91-08, UCLA Comp. and App. Math., 1991.

[7] S. Chandrasekaran and I. Ipsen, *On rank-revealing QR factorisations*, Research Report 880, Department of Computer Science, Yale University, 1991.

[8] J. Charlier, M. Vanbegin, and P. Van Dooren, *On efficient implementations of Kogbetliantz's algorithm for computing the singular value decomposition*, Numer. Math., 52 (1988), pp. 279–300.

[9] J. Delosme and I. Ipsen, *Computing partial correlations from the data matrix*, Research Report 541, Department of Computer Science, Yale University, 1987.

[10] ——, *From Bareiss' algorithm to the stable computation of partial correlations*, Journal of Computational and Applied Mathematics, 27 (1989), pp. 53–91. Also in: *Parallel Algorithms for Numerical Linear Algebra (Advances in Parallel Computing, 1)*, H. van der Vorst and P. van Dooren, eds., North Holland, 1990.

[11] J. Demmel and W. Kahan, *Accurate singular values of bidiagonal matrices*, SIAM J. Sci. Stat. Comput., 11 (1990), pp. 873–912.

[12] K. Fernando and B. Parlett, *Accurate singular values and differential QD algorithms*, Tech. Report PAM-554, Center for Pure and Applied Mathematics, University of California, Berkeley, 1992.

[13] L. Foster, *Rank and null space calculations using matrix decomposition without column interchanges*, Linear Algebra and its Applications, 74 (1986), pp. 47–71.

[14] G. Golub, *Numerical methods for solving linear least squares problems*, Numer. Math., 7 (1965), pp. 206–216.

[15] G. Golub and W. Kahan, *Calculating the singular values and pseudo-inverse of a matrix*, SIAM J. Num. Anal. Ser., B2 (1965), pp. 205–24.

[16] G. Golub, V. Klema, and G. Stewart, *Rank degeneracy and least squares problems*, Tech. Report STAN-CS-76-559, Computer Science Department, Stanford University, 1976.

[17] G. Golub and C. Reinsch, *Singular value decomposition and least squares solutions*, Numer. Math., 14 (1970), pp. 403–20.

[18] G. Golub and C. van Loan, *Matrix Computations*, The Johns Hopkins Press, 1989.

[19] W. Gragg and G. Stewart, *A stable variant of the secant method for solving nonlinear equations*, SIAM J. Numer. Anal., 13 (1976), pp. 889–903.

[20] I. Ipsen, *A close look at an extended Golub-Reinsch algorithm for computing singular values (abstract)*, in The Householder Symposium XI Meeting on Numerical Linear Algebra, Technical Report, Department of Mathematics, Linköping University, Sweden, 1990.

[21] J. Lagarias, *Monotonicity properties of the Toda flow, the QR flow, and subspace iteration*, SIAM J. Matrix Anal. Appl., 12 (1991), pp. 449–62.

[22] B. Parlett, *The Symmetric Eigenvalue Problem*, Prentice Hall, 1980.

[23] B. Parlett and W. Poole, *A geometric theory for the QR, LU and power iterations*, SIAM J. Numer. Anal., 10 (1973), pp. 389–412.

[24] G. STEWART, *Rank degeneracy*, SIAM J. Sci. Stat. Comput., 5 (1984), pp. 403–13.

[25] ——, *An updating algorithm for subspace tracking*, Tech. Report UMIACS-TR-90-86, Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, 1990.

[26] ——, *On an algorithm for refining a rank-revealing URV factorization and a perturbation theorem for singular values*, Tech. Report UMIACS-TR-91-38, Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, 1991.

[27] D. WATKINS, *Understanding the QR algorithm*, SIAM Review, 24 (1982), pp. 427–40.

[28] D. WATKINS AND L. ELSNER, *Self-equivalent flows associated with the singular value decomposition*, SIAM J. Matrix Anal. Appl., 10 (1989), pp. 244–58.

[29] ——, *Convergence of algorithms of decompositions type for the eigenvalue problem*, Linear Algebra and its Applications, 143 (1991), pp. 19–47.

[30] J. WILKINSON, *The Algebraic Eigenvalue Problem*, Oxford University Press, 1965.